POLICY BRIEF: The Sufficiency of Information Protection in AI Governance

To: Artificial Intelligence Safety Institutes & Data Policy Regulators (NIST, FTC, ICO, EDPB)

From: Dr. David Meeler, Specialist in Information Ethics & Legal Philosophy

Date: February 2026

Subject: Moving Beyond "Thick" Ethics: A Practical Reductionist Framework for AI Privacy

1. Executive Summary

Current attempts to regulate Artificial Intelligence are stalling because they focus on "thick" moral concepts (e.g., "fairness," "dignity," or "autonomy"), which are difficult to define legally and nearly impossible to code for technically.

This brief argues for a **"Practical Reductionist"** strategy. Based on foundational research in the logic of privacy, I demonstrate that a conceptually "thin" but rigorously enforced protection of **information** is sufficient to secure the "thicker" moral interests of citizens.

The Policy Bottom Line: We do not need to legislate the complex morality of AI usage. We simply need to secure the information from being *known*. If the information is protected, the moral harms (manipulation, bias, exposure) are preempted.

2. The Regulatory Challenge

Regulators are currently attempting to police the *outcomes* of AI models (e.g., "Does this output violate human dignity?"). This approach is flawed for two reasons:

1. **Subjectivity:** "Dignity" and "Harm" are context-dependent and legally slippery.
2. **Lag:** By the time a "harmful outcome" is detected, the privacy violation (the processing of the data) has already occurred.

The Solution: Shift the regulatory focus entirely to the **Information** layer. A definition of privacy centered on "protecting information from being known" is concise enough to be legally functional, yet robust enough to be morally valuable.

3. Evidence & Analysis: The "Asymmetry of Protection"

My research (*The Monist, 2008*) identifies a critical logical relationship between data and privacy that provides a clear roadmap for AI safety:

A. The Sufficiency of the "Thin" View

Protecting **information privacy** (preventing data from being known/processed) is a *sufficient condition* for protecting "thick" interests.

- *Logic:* If an AI model is technically prevented from "knowing" sensitive attributes (through exclusion or cryptographic blinding), it is logically impossible for that model to use those attributes to manipulate, degrade, or bias against the subject.
- *Benefit:* This allows us to protect complex human rights via simple, binary technical standards (Access: Yes/No).

B. The Failure of the Alternative

Protecting only "access" (physical intrusion) or "expression" (decisional privacy) fails to protect the information itself.

- *Risk:* If we allow the data to be ingested but try to regulate "how it is used," we fail. Once the information is known by the system, the subject loses control, and the "thick" moral harms become inevitable.

4. Recommendations for AI Policy

To implement a **Legally Functional** privacy framework, agencies should adopt the following "Information-Centric" standards:

1. Abandon "Intent-Based" Regulation

Stop asking *why* a company wants to process data. If the information identifies a subject, the protection must be absolute regardless of benevolent intent. The "why" is irrelevant if the "what" (the information) is secured.

2. Enforce "Information Containment" over "Algorithmic Alignment"

Resources currently spent on "Ethics Boards" (debating moral outcomes) should be redirected to **Data Lineage Verification**.

- *Policy Rule:* If a model cannot prove the "Chain of Custody" for a piece of information, it cannot legally hold that information.

3. Codify the "Right to be Unknown"

Move beyond the "Right to be Forgotten" (which implies deletion *after* processing). Establish a "Right to be Unknown": a pre-emptive standard where personal information is legally toxic to unauthorized algorithms by default.

5. Conclusion

We define the parameters of legitimate AI development most effectively not by debating philosophy in the boardroom, but by carefully defining the parameters for protecting

information. A "clean," reductionist information standard is the only shield robust enough to survive the complexity of the algorithmic age.

About the Author

Dr. David Meeler is a Professor of Philosophy specializing in the logic of information privacy, legal theory, and applied ethics. His work focuses on developing privacy conceptions that bridge the gap between "rich" moral requirements and "concise" legal application.

Relevant Research:

- *"Is Information All We Need to Protect?"* (The Monist, 2008) – A defense of informational privacy as the practical core of privacy interests.
- *Dissertation:* "Shared Access & Private Space: A Legal and Philosophical Analysis of Privacy"

Contact: Web: http://faculty.winthrop.edu/meelerd
Sage Profile: https://policyprofiles.sagepub.com/profile/31753/david-meeler